

P Versus Probably P

Introduction:

When one discusses skepticism with undergraduates, many students react to the skeptic's attack on the possibility of knowing some proposition P by retreating to the safer claim that one can at least know that probably P. If one responds this way to the worry posed by traditional skeptical scenarios, consistency might seem to force one to the odd sounding conclusion that there are very few contingent propositions one should claim to know. One should, instead, only claim to know that such propositions are probably true. While many are happy to move in that direction, they start to get nervous if they become convinced that the justification they have for believing propositions about probability is itself *fallible* justification, i.e. justification that is, as far as they can tell, consistent with the proposition about probability *itself* being false.¹ They don't want to be backed into a corner in which, if they are to be precise, they should concede of a contingent proposition P that they only know that it is probable that P is probably true. And once they have admitted fallibility with respect to first-order probability claims, of course, it's hard to see how they can suddenly become infallible with respect to second-order claims about what is likely to be likely. And now we are in the unfortunate position of being unable even to complete our knowledge claims. We know only that it is probable, that it is probable, that it is probable....., and so on *ad infinitum* as we search in vain for a way to complete our assertion.

Hume (1888, 184-85) thought that there was an argument for skepticism in the neighborhood of the above observations. But the argument he offered was problematic. Hume suggested that if P is only probable for you, then it has a probability of less than 1—let's stipulate that the probability is .9. He then argued that a healthy fallibilism should lead one to concede that it is only probable that the proposition has that probability of .9, and if that is so, its real probability is less than .9, say .8. But,, and we repeat the argument until we reach the conclusion that the proposition in question has a probability as close to 0 as it can be without being 0. The error in the argument, of course, is that when it comes to judgments about probability, one can miss one's mark higher or *lower*. So if I am fallible with respect to the original judgment that P has a probability of .9, I *might* have *underestimated* its probability. There is then no reason for me to let my realization that I am fallible force me to *lower* the original estimate of probability.

Even if the above argument for skepticism is a bad argument, it would still be disastrous if fallibility forces one to the problematic view that one can never reach a conclusion about the probability of a given proposition. As one tries to formulate what one is permitted to believe, one will have an infinitely complex proposition involving infinitely many iterations of "it is probable that..." That is arguably worse than some radical skeptical conclusion—so bad that we should surely think twice about admitting that if our evidence only makes probable a given

¹ Defining infallible justification is much harder than one might think. Things start to go badly once one notices that every necessary truth is entailed by all of one's evidence. It is a long story, but I think one needs a relevant logic to solve all the difficulties one encounters.

proposition P, then all that we should claim to know, or even justifiably believe, is that it is probable that P. This seems right to me, but in what follows I'll explore reasons, some older, some more recent, for thinking that one should protect what one claims to know and justifiably believe with at least one round of probability operator.

The Abominable Conjunction:

Many would argue that it is simply a datum of linguistic usage that one cannot without “dissonance” make the following sort of claim: “I know that P but it might not be true.” Almost as many hear the same sort of “dissonance” if someone simply asserts: P but it might not be true. Suppose you are at a trial and the prosecutor asks: Do you know that the man before you is the man you saw leaving the room? You are going to confuse everyone if you answer: Yes, I know that was the man, though of course it might not have been. They are probably going to be as confused if the prosecutor simply asks: Is that the man you saw leaving the room, and you respond: It is the man who left the room, but it might not be. It may be that Williamson and his many followers can explain this last datum by convincing us that “knowledge is the norm of assertion.” And even if it isn't, in general, the norm of assertion, it certainly might be when one is testifying in a courtroom. So when on the witness stand you identify the man who left the room, your “audience” is still going to *infer* that you are committed to the view that you know which man left the room. You didn't *assert* that, but both you and your audience understand that it would have been inappropriate for you to make that assertion in that context without the relevant knowledge.²

The most obvious explanation of why you can't legitimately claim to know that P while you admit that not-P is possible is that that the possibility in question is most naturally construed as *epistemic* possibility.³ And the most natural understanding of claims about epistemic possibility is the following: Not-P is epistemically possible for S if not-P is logically consistent with the conjunction of everything S knows.⁴ But if S knows that P then, trivially, not-P is inconsistent with the conjunction of everything S knows.

There are difference responses to the abominable conjunction. One is to rethink more carefully one's willingness to allow that certain propositions are epistemically possible. The eyewitness who saw Fred leave the room does, indeed, know that Fred left the room, and

² This distinction between what one asserts and what conventions that govern assertions also explains Moore's paradox. It is decidedly odd to assert a sentence of the form “P but I don't believe that P.” The sentence is not unintelligible—indeed it can be true. But in asserting the first conjunct one implies to one's audience that one does believe P, something that one then denies with one's second conjunct.

³ If you were interpreted as asserting only that you know P while P is, for example, a contingent proposition (and is possible in that sense), there is no puzzle.

⁴ This isn't quite right. Every necessary falsehood is incompatible with whatever people know, but there is surely a sense in which some (complex) necessary falsehoods are epistemically possible for me. It is not clear how to solve this problem without introducing something like a relevance logic, or an awareness requirement of some kind. With the former, we discount the significance of irrelevant inconsistency. Necessary falsehoods are inconsistent with everything we know and believe and we aren't going to count that sort of inconsistency as revealing relativized epistemic impossibility. On the latter view, we discount the relevance of inconsistency when the subject for whom P is supposed to be epistemically impossible is unaware of the inconsistency.

shouldn't admit that it is possible (relative to the evidence) that it was anyone else. This is a tough row to hoe, however. When the prosecutor brings up skeptical scenarios, albeit typically scenarios less fantastic than the philosopher's favorite examples (brains in a vat, malicious, very powerful deceivers, and the like), it strikes most people that it is almost silly to deny the relevant possibilities. *Couldn't* Fred have a long lost identical twin brother? *Couldn't* Fred have a doppelganger? *Couldn't* you have been under the influence of mind-altering drugs that you ingested without your knowledge? It seems almost obvious to most people that it would be inappropriate to deny the respective possibilities, even when those possibilities are construed as epistemic. In light of these observations, one might conclude that if we were being precise, it would be wise to retreat to *probability* claims, both with respect to what we know and what we assert. So the careful witness could assert without air of paradox that he knows that it is almost certain that he saw Fred leave the room, while also admitting that, of course, there is an outside chance it wasn't. Or without talking about knowledge, the witness could just assert it was very likely, almost certain, that he saw Fred leave the room (though of course there is an outside chance it wasn't Fred).

There are a number of suggestions that have been proposed to accommodate the above rather robust linguistic data, some of which do not turn to probability. There are, for example, many different versions of contextualism (Cohen 2001, DeRose (1992), Lewis (1996,xx,xx)). The crude idea behind at least one version of the view is that the truth conditions for knowledge claims, and more generally epistemic claims, vary depending on context. And one of the things that can change the relevant context, according to some (see Lewis 1996), is the fact that the person making a knowledge claim is specifically asked to consider a certain skeptical scenario.⁵ On one such view, once a possibility is raised, you need to have evidence to eliminate it if you are to continue claiming that you know a proposition that is inconsistent with that hypothesis. So it was OK for the witness to claim to know it was Fred who left the room *until* Fred's defense lawyer started raising the possibility of a long-lost twin.⁶ But once that scenario is raised, the witness needs evidence strong enough to eliminate that possibility. And typically the witness doesn't have such evidence—not evidence that *eliminates* the possibility.

The literature on contextualism is voluminous, and I don't intend to survey it here. It certainly would be odd if a defense lawyer could wring this sort of concession out of a witness merely by raising a possibility. Were this the case, those charged with serious crimes, particularly those who are guilty, should probably hire for their defense not lawyers, but epistemologists steeped in the history of skeptical scenarios.

Another approach, distinct from, though consistent with contextualism, is to deny closure principles. The argument sketched above seems to rely on a crude closure principle⁷ of the following form: If you really do know that P and P is incompatible with Q then you are at least

⁵ Other versions of the view stress that what is at stake if one is wrong is critical to changing the context.

⁶ The criminal justice system knows that it had better place limitations on the introduction of "skeptical scenarios." That's why it will require of the defense lawyer trying to introduce some possibility a "foundation"—some reason for thinking that it has a reasonable chance of being true.

⁷ Here we are just talking about so-called single premise closure. As we'll see below, multi-premise closure is far less plausible and presents problems of a different sort.

in a position to know that not-Q based on P. If our witness knows that Fred is the one he saw leave the room and Fred's being the one who left the room is incompatible with its being his doppelganger, Frodo, who left the room then the witness should be in a position to know that it wasn't the doppelganger who left the room. To the extent that the witness is uncomfortable making such a claim, that's a reason for the witness to rethink the original knowledge claim. Again one can deny closure without embracing any particular account of knowledge. But the contextualist is there to explain why closure seems to fail as the context of a conversation changes and relevant alternatives come in to play.

The contextualist isn't the only epistemologist who rejects even simple closure principles of the sort sketched above. Nozick famously endorses a sensitivity condition on knowledge. You only know that P only if (among other conditions) you wouldn't believe that P were it the case that not-P.⁸ Nozick doesn't provide an analysis of the truth conditions for subjunctive conditionals, but does appeal to the familiar metaphor of possible worlds. If you want to know whether you would not believe P were P false, take yourself "imaginatively" to the "nearest possible worlds" in which not-P is the case, and ask yourself whether in such worlds you would still believe that P. In the case of mundane beliefs, like my belief that I am in my office now, the idea is supposed to be that I wouldn't believe that I was in my office if I weren't. At least that's so if the world is as I think it is. After all, the "close" possible worlds in which I am not in my office are worlds in which I decided to do something else. And if I were, for example, on the golf course, I wouldn't believe that I was in my office. On the other hand, most skeptical scenarios are described in such a way that a) I believe they do not obtain, and b) I would still believe that they didn't obtain, even if they were to obtain. So I don't believe that I am a brain in a vat manipulated so as to have the very experiences I am having right now, and I would believe precisely that even if I *were* a brain in a vat. A moment's reflection tells us that if Nozick's view were correct—more generally if there were a sensitivity condition on knowledge—we should not endorse closure.

Again, I won't evaluate Nozick's view here. I have done so elsewhere (xx). Suffice it to say, that at least some take the denial of single premise closure to be a *reductio* of the view. Both contextualism and sensitivity conditions on knowledge are painful enough remedies for our problems, that it would be worth exploring alternatives, and some of those alternatives will pull us again toward protecting our claims with probability operators.

Descartes is usually interpreted as a philosopher whose quest for secure knowledge led him to employ his famous "method of doubt." To test a candidate for genuine knowledge (or at least genuine foundational knowledge), ask yourself whether the proposition believed could be false given your epistemic situation. Because Descartes thought that massive deception with respect to the physical world and (he seemed to argue) even elementary mathematical propositions was intelligible, he *initially* ruled these out as candidates for knowledge. Famously,

⁸ It gets more complicated. Nozick modifies the view to make reference to methods employed in reaching a conclusion. If you know that P employing some method of believe acquisition M, then you wouldn't believe P using M were P false. Nozick also endorses a distinct safety condition: One's belief that P satisfies the safety condition of one not only has a true belief that P but in close possible worlds one also believes truly that P. We'll say more about safety shortly.

he hit upon his own existence as a truth that survives the test for what he could know, and somehow (this gets much more mysterious) he thought he could learn enough from that knowledge to also know various premises (including one about God's existence) that would negate his initial grounds for doubting more commonplace propositions about the world around us. (Cunning xx, clarity and distinctness xx)

Descartes's idea that knowledge requires being in a position to eliminate the possibility of error obviously sounds remarkably similar to Lewis's assertion that one knows that P only if one can eliminate the possibility of error. The difference is supposed to lie with Lewis's view that the alternatives in question are only those that are relevant (though, on his view, one had better not *mention* far-fetched alternatives lest for that reason they become relevant). There is a sense in which certain forms of contextualism really aren't that far removed from Descartes's criteria for foundational knowledge.

The Cartesian conception of knowledge, like the contextualist's account of knowledge, both get support from the ease with which we can get most people to retract knowledge claims. I suppose it depends on how one teaches the course, but it has always been my experience that students in *Introduction to Philosophy* barely need an argument to retract knowledge claims. As soon as one gets a dream argument on the table, many students shrug shoulders and concede that one can't *know* that one isn't being deceived in one way or another about the physical world, the past, other minds, or the future. A lot depends on the inflection with which one uses the word "know"—the more emphasis one places on the word, the more likely it is that the relevant concession will be made.

Lottery Puzzles:

If one doesn't concede the knowledge skepticism discussed above, the skeptic might appeal to more mundane data concerning how we use the expression "know." Consider the much discussed lottery puzzle and preface "paradox." If anything is obvious it is that it is very unlikely that our recently bought "Lotto" ticket will be a winner. You probably have a better chance of having a heart attack in the next few hours than you have of winning the lottery with the purchase of a single ticket. And it is not unusual for an author to concede in a preface to a book, that the work probably includes errors--this despite the fact that we may suppose that the author is relatively confident of each individual assertion made in the book. Few philosophers adopt Cartesian criteria for knowledge. One of the most common complaints is that such a view renders false almost all of our ordinary claims to have knowledge. Do you know what time it is? Do you know who the first president of the U.S. was? Do you know where the Declaration of Independence was signed? Do you know what you had for breakfast? Outside of a philosophical discussion, all sorts of people will answer these questions, "Yes." There is surely something wrong with a view that renders false commonplace assertions with which most people are perfectly comfortable.

For a while many epistemologists advanced as an alternative to the idea that knowledge requires infallible justification, the view that knowledge is better understood as justified true belief. This account of knowledge invites at least two questions: 1) How much justification is

necessary for knowledge? and 2) How strong a belief is necessary for knowledge? And if we allow that one can know some proposition P without having justification that guarantees the truth of P, one faces all sorts of Gettier counterexamples. Still, to this day, a great many epistemologists hold out hope that we can find conditions to add to justification, belief and truth that will mend the analysis and preserve the heart of the idea behind knowledge as justified true belief. But however one tries to mend the analysis the fact remains that almost no one is inclined to claim that we know of a given lottery ticket that it is a loser. Still, the justification we have for believing that the ticket is a loser is incredibly strong. In other contexts we don't hesitate to claim knowledge of this or that with far less evidence than we have for thinking of each ticket that it is loser. When pressed as to why it is inappropriate for someone to claim knowledge that the ticket just bought is a loser, almost everyone will point out that there is at least a chance—no matter how small a chance—that the ticket will be the winner. And the implication is that the existence of that chance is incompatible with insisting that one knows that the ticket will not win. The strongest knowledge claim one should make is that one knows of each ticket that it will probably lose.

Parity of reasoning might suggest that if our justification falls short of what is required for knowledge in the case of our justified belief that a given lottery ticket is a loser, it falls short of what is required for knowledge in a host of other cases. To consider an often used example, I might ordinarily claim to know where my car is—I remember parking it my office parking lot. But when pressed, I will quickly acknowledge that there is at least a chance that someone has stolen the car from the lot. I might ordinarily claim to know where I was born. But when pressed, I will probably acknowledge that there is at least a chance that my parents fabricated the story they told me about my birthplace. I might ordinarily claim to know what I had for breakfast, but when pressed I might concede that there is at least a chance that I am confusing what I had for breakfast yesterday with what I had for breakfast today. If the chance that the lottery ticket is a winner precludes my legitimately claiming to know that it is a loser, why don't all these other "chances" preclude the legitimacy of my claiming to know that which is incompatible with the relevant possibilities?

In response to the argument from parity one might hope that there are conditions for knowledge unrelated to degree of justification, conditions that are met by many mundane beliefs and not met by a belief that a given lottery ticket is a loser. Earlier we looked at sensitivity. Even if the ticket I hold in my hand were a winner, I still would believe, based on my statistical reasoning, that it is a loser. By contrast, it is not clear that if I were born someplace other than Toronto, I would still believe that I was born in Toronto.⁹ As we say, however, a condition of sensitivity on knowledge is highly controversial--it will almost certainly lead us to deny very plausible claims concerning knowledge closure. Others suggest that we substitute for a sensitivity condition a safety condition. One doesn't know that P if there are close possible worlds in which one would believe P falsely.¹⁰ And some would argue that if I were to believe

⁹ As we discussed earlier, the "close" possible worlds in which I was not born in Toronto are worlds in which I was born in a nearby city, and they are also worlds in which I know that. As we'll see later I'm not at all sure that there is any real explanatory value to talk possible worlds that are near and remote.

¹⁰ A safety condition might help with Gettier counterexamples like fake barn country.

that the lottery ticket I just bought will lose, the belief would fail to satisfy a safety condition. There are supposed to be close possible worlds in which I believe *falsely* that the ticket in my hand is a loser. Frankly, however, I have never understood the argument. Are there close possible worlds in which I believe falsely that the ticket I'm holding will lose? I suppose I can make sense of a world in which we would be tempted to say such a thing. Suppose that as it turns out the winning ticket is numbered 259887324 and the ticket I just bought is numbered 259887325. Upon discovering the number of the winning ticket I might well exclaim "I was so close!!!" If I were philosophically educated I suppose I might even say that there was an extraordinarily close possible world in which I would have held a winner. But there are as many possible worlds as there are numbered tickets. While it is true that each world in which the winning ticket is not mine is "equidistant" from the actual world, the fact is that all of those worlds are *remote* possible worlds.

The problem is that we are misled by the metaphor of close and remote possible worlds. We need to keep reminding ourselves that it is a *metaphor*. There is only one world and everything that is true (and that includes counterfactuals, and truths about probability, necessity and possibility) is made true by virtue of some feature of the one and only actual world. The worry is that philosophers sometimes appeal to the metaphor of near and distant possible worlds to avoid looking long and hard for actual truth makers. When we say, for example, that the world in which I am a brain in a vat is remote (far removed from the actual world), I suspect that we mean nothing but that our evidence makes it very unlikely that we are brains in a vat. And that might be true. But if we are epistemologists our task is to explain what the available evidence is and how it makes probable or improbable various hypotheses. The world in which the lottery ticket I just bought will win is remote—and that's precisely because relative to my evidence it is extraordinarily unlikely that the lottery ticket will win.

It's not clear to me that one gains much by way of bolstering the above argument with more technical philosophical distinctions. Nevertheless, earlier we briefly discussed the idea of closure. If you really do know that P and you also know that P entails Q, then, at the very least you are in a position to know Q. The principle as stated is sometimes called the principle of single-premise closure. But multi-premise closure principles for knowledge strike many as almost as plausible. If you know that P1, you know that P2, you know that P3..., and you know that Pn, while you also know that the conjunction of P1, P2, P3..., Pn entails Q then you are at least in a position to know Q. But multi-premise closure spells disaster for anyone willing to claim that we do, in fact, know of each ticket in the lottery that it is a loser. That position, together with multi-premise knowledge closure, entails that we know that the lottery won't have a winner. And that really is an absurd conclusion.

The situation isn't nearly as dire for the philosopher who insists that we have a justified belief of each ticket that it is a loser. And that's because there isn't much to be said for multi-premise closure principles governing justified belief. It just seems wrong to suppose that if you are fallibly justified in believing P1, P2, P3, P4..., and P100, then you are justified in believing the conjunction of these propositions. Indeed, you are surely justified in believing that the conjunction is false (assuming that the propositions are logically and probabilistically

independent). That's how airlines book rationally book tickets. That's why we don't think that the best NBA player at making free throws will make all of his free throws in the course of a season. So, one might argue, there is no reason to suggest that we are only justified in believing of each ticket that it will probably lose. But here the difference might be merely semantic. On at least one natural interpretation, to say that S is justified in believing P just is to say that P is epistemically probable relative to S's epistemic situation. To insist that we are only justified in believing that a given proposition P is probable is to move up a level and claim that it is only probable that P is probable. That is an *intelligible* assertion, and it may be an assertion one wants to make, but as we said in our introductory remarks, one is in danger of committing oneself to the view that we can only make claims protected by infinitely many iterations of probability qualifications—the view that it is only probable for S that it is probable for S that is probable for S..., and so on *ad infinitum*. And one would need independent reasons for supposing that one is forced to such a conclusion.

If one were a Cartesian about knowledge, if one embraces the view that knowledge is the norm of assertion, if one thinks that most ordinary knowledge claims fail to satisfy Cartesian criteria for knowledge, and if one thinks that even most claims to know what is likely (relative to one's evidence) also fail to satisfy Cartesian criteria for knowledge, *then* one might be forced to the strange view that one should not even assert of a given lottery ticket that it is likely to be a loser. One should only assert that it is likely that it is likely that it is likely, and so on *ad infinitum*, that it is a loser. And if talk about S's justified belief that P can be translated into P's being probable relative to S's evidence, one could also reach the conclusion that one can only legitimately claim to be justified in believing that one is justified in believing that one is justified in believing, and so on *ad infinitum* that P. But those were a lot of "if"s. It is not at all obvious that we will be forced to this highly counterintuitive conclusion.

It is certainly hard to meet Cartesian criteria for knowing an ordinary proposition describing the physical world, the past, the future, other minds, and theoretical entities posited by science. But is surely easier, at the very least *safer*, to claim knowledge that certain propositions are probable relative to one's evidence. That will depend on how one construes claims about the probability of a proposition relative to evidence. If one's evidence E's making probable P were a contingent matter, knowable only *a posteriori*, it is hard to see how the probability claim would be much safer than a claim about P. If probability has anything to do with the frequency with which a proposition like P is true relative to propositions like E's being true, the probability claim is a decidedly ambitious claim about statistical generalizations—a claim that ranges over past, present, and future. The heart of the reliabilist's concept of inferentially justified belief, just is this idea that the probability of a given proposition is a function of the frequency with which certain sorts of output beliefs are true given that they are caused by certain sorts of true input beliefs. When the claim about frequency/reliability is contingent, one will presumably need to employ some sort of inductive reasoning to justifiably believe that the claim is true.

I have long argued that the best hope for attaining philosophically satisfying justification is a Keynesian notion of epistemic probability. On such a view the proposition that E makes

epistemically probable P, when true, is necessarily true. And the proposition is knowable *a priori*. To be knowable *a priori* isn't to be known *a priori*, and one needs to do a great deal of work to make plausible the view that *a priori* justification is a potential source of infallible justification. But if one can make good these promissory notes, one might be in a position to stop the pressure to iterate probability operators. While I might not be able to have Cartesian knowledge that the lottery ticket I just bought is a loser, I might be able to have Cartesian knowledge that the ticket I just bought is likely (relative to my evidence) to be a loser. What is true of beliefs about lottery tickets might be true of all sort of other relatively mundane beliefs we have. To be sure, when asked if we know where someone is born we just answer "Yes," if we have a strongly justified belief. It would be decidedly odd to answer this question: "No, but I know where the person was probably born." But while that is so, we have already had occasion to note also that many people are quite willing to retract knowledge claims in the face of even mundane skeptical worries. The contextualist has one explanation of this phenomenon, but there is another equally plausible explanation. We just aren't careful and we realize that we just aren't careful in what we casually say. For most purposes, we don't need to be worry about being careful. But when someone wants to make an issue over whether an epistemic claim is accurate we typically realize that we should simply concede the objection.

I'm not suggesting that a retreat to probability is the only way that we recalibrate our epistemic claims. I have argued elsewhere (xx) that it is equally likely that when challenged as to whether we really know what we claimed to know, we will often retreat to conditionals *together* with presuppositions about what is likely. When I say I know where I'll be spending the summer and someone points out that I might be dead long before summer arrives, it is entirely natural for me to admit that all I really know is where I *plan* to be. And if things go as I expect they will, I'll be where I plan to be. I'm not suggesting that the truth conditions for "S knows that P" are best understood in terms of what would be required for S to know some conditional of the form: If such and such, then P. For one thing there is a robust intuition that one can't know what it is false. If I claim to know that P and P is false, my self-attribution of knowledge is false. And that seems to be true even if the person making the knowledge claim is willing to "retreat" in the face of opposition to a more modest claim about intentions or conditionals.

Evidence and the Possibility of Justified False Belief:

There are more recent views about the nature of evidence and the nature of justification that *might* suggest an important role for probability judgements to play. I'm not convinced that either of the views I discuss below is true, but they are important to consider.

Timothy Williamson (2000) has famously endorsed the view that the concept of knowledge is unanalyzable and that our evidence should be identified with what we know. On the face of it such a view allows for there to be justified false beliefs. But a justified false belief cannot itself be evidence for some other proposition if we identify evidence with knowledge.

Such a view is at best controversial. Suppose, for example, that I hear a knock on my office door. As I entered my office a few minutes earlier I noticed that Diane was only colleague at her office. Naturally enough I infer that it is Diane who is knocking at my door. I'd also naturally infer that Diane wants to talk to me. As it turns out, the person at my door is, instead, a stranger who just happened by and was looking for directions. So what should we say about all of this. Let's concede that I know that there is a knock at the door. Let's also concede that I am justified in believing that it is Diane, and am also justified in believing that Diane wants to talk to me. Still, I know neither that it is Diane nor that Diane wants to talk to me. So what is my evidence for thinking that Diane wants to talk to me. One might suppose that one can simply go all of the way back to the knock on the door, but it's not clear to me that one can dispense with the intermediate premise in order to get to the conclusion that Diane wants to talk to me. At the very least my *doxastic* justification might require the intermediate premise (assuming that I base my belief on it and wouldn't have believed that Diane wants to talk to me but for my believing that it is she who is knocking at the door).

If one is convinced that evidence is knowledge and that a justified belief must be based on evidence one is left searching for a proposition upon which to base the relevant belief. Probability can come to the rescue. While I don't know that Diane is knocking at the door, I might know that it is likely (relative to my evidence) that Diane is knocking at the door. And it is not obvious that the latter proposition is false. Why can't that truth about what is likely serve as evidence for the proposition that Diane wants to talk to me.

Littlejohn (2012 Chapters 3 and 4)) takes Williamsonian intuitions further and concludes that it actually makes no sense to suppose that there are justified false beliefs. His view is related to the controversial idea that evidence is factive. Whether a belief is noninferentially justified or inferentially justified, he argues, it must be based on some fact. When I'm noninferentially justified in believing that I'm in pain, my belief is based on the fact that I'm in pain. The argument relies for much of its force on the idea that if S has a justified belief S must hold that belief for the right reason—there must be a true statement of the form S believes that P because of E, where E is a factual claim about the world. If noninferentially justified beliefs are beliefs caused by their truthmakers then it follows directly that there can't be any noninferentially justified false beliefs. If one's inferentially justified belief is one validly deduced from a noninferentially justified belief, then there can't be any inferentially justified false beliefs of *this* sort either.

So far Littlejohn might have some very traditional epistemologists on board. But we haven't considered yet the most obvious candidates for inferentially justified false beliefs—beliefs formed as a result of legitimate, but not deductively valid inference. Here Littlejohn turns to another principle that he takes to be relatively unproblematic. He calls it xxx: If S is justified in believing P and there is at least one really obvious logical consequence of P (for S), say, Q, then S could justifiably believe Q on the basis of P. But if Q could be justified *because* of P then P would need to be true. "Because of" is factive. So if P is really a justified belief, one on the basis of which one can believe something else, then P would need to be true.

Littlejohn relies on a number of controversial premises. And those premises had better be highly plausible given the counterintuitive nature of his conclusion. It seems to most almost obvious that one can have a justified but false belief. Let's set aside any radical skeptical concerns for the moment. If I meet you at a party and you tell me that your name is "Tom," I am surely justified in believing that "Tom" is your name. And that is true even if you are a pathological liar who takes perverse pleasure in misleading people with respect to your name. Williamson tries to soften the counterintuitive nature of the position by relying on a distinction that sometimes seems to be made in the law between behavior that is legally impermissible despite the fact that we would excuse the person for acting in that way. So consider the police officer who mistakes a person holding a toy gun for a person holding a real gun and kills that person believing that she was acting in self-defense. On one view, the police officer acted in a way that was not legally justified. There was no legitimate claim of self-defense. On the other hand, most of us would excuse the behavior—we would think it wrong to punish, or even blame the police officer for acting as she did. Similarly, one might argue, I was unjustified in thinking that your name was "Tom" even though it would be inappropriate for anyone to blame me for believing what I did. My doxastic attitude was perfectly *understandable*. I would have a perfect *excuse* for believing what I did.

The analogy has never seemed to me helpful. But one of the reasons is that I have never been persuaded that the distinction in law is itself intelligible. Let's focus again on self-defense. It seems to me that one legitimately kills in self-defense when one kills with a reasonable belief that one's life is threatened. If you approach me with a realistic toy gun pointed at my head, and I have the opportunity to eliminate the rationally perceived threat you pose, then, it seems to me, I have acted in a way that is perfectly legitimate. There are civil laws framed in terms of strict liability. For various reasons we sometimes hold people responsible for harms they cause even if those people had no reasonable belief that the harms in question would result from their actions. And, of course, one could have analogous "strict liability" with respect to the way in which we frame certain criminal laws. But morality is different from legality, and even if we make someone legally responsible for killing someone who, in fact, posed no actual threat, it wouldn't alter the fact that the person who killed with a reasonable fear for his/her life was acting in a moral/rational way. And the relevant analogy for epistemology isn't what the law requires, but is, instead, what morality/rationality requires.

In this context, however, I only want to suggest that if one is tempted by the idea that one cannot have a justified but false belief, probability claims are a much more plausible way of softening the counterintuitive consequences of such a view. If I can't be justified in believing that Diane knocked on my door, and that Diane wanted to talk to me, I can surely be justified in believing that it is likely relative to my evidence that Diane knocked on my door and that Diane wanted to talk to me. Those propositions are true. And from the fact that it is likely that is Diane knocking on my door, I can legitimately infer that it is likely that Diane wants to talk to me. Even if we grant Littlejohn his premises, the above claim presents no problems.

References:

Cohen, Stewart, 1986. *Knowledge and Content*. Journal of Philosophy 83: 574-583

De Rose, Keith. 1992. Contextualism and Knowledge Attributions. Philosophy and Phenomenological Research, 52, 913-929.

Hume, David, 1888. *A Treatise of Human Nature*. Ed. L. A. Selby-Bigge. London: Oxford University Press.

Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of philosophy* 74: 549-67.

Littlejohn, Clayton. 2012. *Justification and the Truth-Connection* (Cambridge University Press).

Fantl, Jeremy and Matthew McGrath. 2009. *Knowledge in an Uncertain World*. London: Oxford University Press.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.